Streaming Readout for EIC Detectors

Proposal submitted 25 May, 2018

STREAMING READOUT CONSORTIUM

S. Ali, V. Berdnikov, T. Horn, I. Pegg, R. Trotta
Catholic University of America, Washington DC, USA
M. Battaglieri (Co-PI)[1], A. Celentano
INFN, Genova, Italy
J.C. Bernauer* (Co-PI)[2], D.K. Hasell, R. Milner
Massachusetts Institute of Technology, Cambridge, MA
C. Cuevas, M. Diefenthaler, R. Ent, G. Heyes, B. Raydo, R. Yoshida
Thomas Jefferson National Accelerator Facility, Newport News, VA

* Also Stony Brook University, Stony Brook, NY

ABSTRACT

Micro-electronics and computing technologies have made order-of-magnitude advances in the last decades. Many existing NP and HEP experiments are taking advantage of these developments by upgrading their existing triggered data acquisitions to a streaming readout model. A detector for the future Electron-Ion Collider will be one of the few major collider detectors to be built from scratch in the 21st century. A truly modern EIC detector, designed from ground-up for streaming readout, promises to further improve the efficiency and speed of the scientific work-flow and enable measurements not possible with traditional schemes. Streaming readout, however, can impose limitations on the characteristics of the sensors and sub-detectors. Therefore, it is necessary to understand these implications before a serious design effort for EIC detectors can be made. We propose to begin to evaluate and quantify the parameters for a variety of streaming-readout implementations and their implications for sub-detectors by using on-going work on streaming-readout, as well as by constructing a few targeted prototypes particularly suited for the EIC environment.

[1] battaglieri@ge.infn.it
[2] bernauer@mit.edu

PROPOSAL

Introduction

The Electron-Ion Collider is proposed as a facility to probe the fundamental structure of matter through lepton scattering with unprecedented luminosity from the nucleon and nuclei. The collider topology offers great advantages in the ability to detect both the current and target fragmentation regions, i.e., the detector is expected to be able to approach complete reconstruction of the hadronic final-state. The EIC detector will take data at high collision luminosity in a fixed configuration for extensive periods. Traditional triggered data acquisition from such a detector raises the possibility that certain observables, not known to be important at the time of data-taking, will be missed. Streaming readout has the potential to ensure that the most complete set of measurements with a given detector configuration will be recorded and be available indefinitely for analysis. Given the tremendous cost and effort to realize EIC and its scientific equipment, it seems prudent to have a modest R&D effort at this time to study implementation of streaming readout.

In it's simplest form, streaming readout is the continuous collection of data from the detectors without any selection by a hardware trigger. However, for EIC scale detectors, simply writing out all data would very likely exceed the ability to write to and keep them on storage. A functional streaming readout design works in multiple stages of data reduction, from per-channel zero-suppression already found in standard electronics, to the use of high-level analysis involving sophisticated processes like track reconstruction. Many of these processes will require on-the-fly calibration, extracted from the data during data taking itself. In streaming-readout schemes, these selection algorithms will be written in high-level languages–rather than designed in hardware, or programmed in FPGA-firmware–making them more accessible and easier to verify. This data reduction allows, ideally, the recording of all relevant data while rejecting uninteresting information. The accessibility of full detector information online will also vastly improve the experiment's monitoring capability

Some current generation experiments were designed in the conventional triggering scheme and evolved into streaming readout as technology advanced. For the Electron-Ion Collider, given the tremendous improvements in micro-electronics speed, in network bandwidth, and advanced computing capabilities, streaming readout can and should be planned from the beginning with the aim to maximize efficiency and flexibility of the EIC compared to a conventional approach. However, such innovation will not only impact the readout scheme and offline processing, it will imply certain choices for detector technologies to be deployed. It is therefore necessary to begin to understand the parameters of streaming-readout technology now so that reliable extrapolation to the final parameters, that can directly determine major choices in detector design, can be made.

LHCb is an example of an experiment that has recently deployed streaming readout. This development has enabled the collaboration to decrease the time-to-publication from months-to-years, down to weeks. It is interesting to note that the streaming readout requirement at LHCb has already had repercussions on their detector upgrade plans. For example, pixel detectors were chosen for the upgrade over strip detectors in order to speed up track-finding in real-time.

Since the decision to employ streaming readout framework for EIC will have significant impact on the sub-detectors design, the availability of a robust framework at the time of design, prototyping, and testing of detectors will benefit these activities immensely. Understanding the implications of streaming readout will allow for a better standardization of the system, making individual ad-hoc solutions unnecessary and avoiding costly retro-fitting.

For this reason, the R&D program on streaming readout, with focus on the performance improvement with respect to traditional solutions, should start as soon as possible as part of the "generic" R&D program. This cannot be postponed to pre-engineering studies, when the overall design will be already defined. The development of a streaming readout framework will be relevant to all detectors and all potential EIC sites.

R&D for EIC

In this phase of R&D, we aim to identify and quantify relevant parameters in evaluating streaming readout at the EIC by analyzing prototype implementations of streaming readouts. We will partially make use of prototypes streaming-readouts that are already being developed with applications other than EIC in mind. In addition, we propose start a few small-scale prototype streaming-readout projects to evaluate some specific aspects for EIC detectors.

Deliverable is a document/publication containing:

- Reports on the relevant aspects of the performance of the prototypes studied.
- A list and definition of streaming-readout parameters relevant to the EIC.
- Initial estimates of some of the parameters with current technology as well as extrapolations to the time period of EIC detector construction phase.
- A list and definitions of relevant parameters for detector technologies (e.g. TPC, Crystal Calorimeter, etc.) when considering streaming readout.
- Initial estimates of some of these parameters.

With this deliverable in mind, we have identified several on-going, or independent prototype projects to be a part of this effort. **There will be no funding requests for the following projects**

TPC in streaming-readout at JLab A Tagged Deep Inelastic Scattering experiment, TDIS, has been proposed for hall-A at Jefferson Lab. TDIS measures coincidence between scattered electrons in the Super BigBite spectrometer and low momentum recoil protons in a GEM based radial time projection chamber (rTPC). The rTPC detector is an ideal candidate for streaming readout. A rTPC readout test stand is being put together at Jefferson Lab. using components used in the ALICE data acquisition upgrade. Procurement of components and access to the required software and firmware is underway with a target date of July 1st, 2018 for the test stand to be operational. One of the goals of this project is to gain experience with rTPC readout in a streaming mode. This experience will be used to guide the design and implementation of a streaming readout for the TDIS rTPC. It will also provide valuable input into the planning process for a generic streaming readout architecture to be used by near future Jefferson Lab. experiments and, potentially, with EIC era detectors.

Crateless-Streaming at JLab The JLab Fast Electronics and Data Acquisition groups are investigating alternatives to VME based readout towards high-rate streaming solutions, taking advantage of existing electronics. A prototype new readout solution using existing JLab-designed front-end electronics, the 250 MHz Flash ADCs (FADC250) and the Drift Chamber Readout Boards (DCRB), will be used as VXS trigger path for readout data instead of the VME bus. The JLab-designed VXS Trigger Processor (VTP) is connected to each front-end module using the VXS backplane and has a dedicated bandwidth of up to 20Gbps to each FADC250 and up to 10Gbps to each DCRB. The VTP also has significant buffering (4GBytes with up to 200Gbps bandwidth) and a 40Gbps optical Ethernet port that allow for a significant improvement in readout data rates per crate compared to using the 200MB/s VME bus that is currently shared for all modules.

We further plan to prototype a streaming option using components of this new readout solution. By buffering large time windows of continuous data (50us) and zero suppressing data we can create a streaming readout system that supports dead-time. Dead-time can be created when front-end input buffers are too full - either because of continuous high occupancy conditions and/or network back pressure. When dead-time occurs all modules in the system will skip sending data for the next time time window and allows front-end module buffers and/or network traffic to empty which eventually allows data taking to resume. Using reasonable time windows and buffering this style of streaming data acquisition system allows the network bandwidth to be defined by the occupancy of the detector (rather than the bandwidth needed for continuous 100% occupancy) and also eliminates efficiency losses related to insufficient network bandwidth where data would be dropped, towards little-to-no dead-time in the streaming. Interface protocols will be used between modules that are not tied to VME, but rather to use Ethernet compatible frames and high level protocols (i.e. UDP, TCP), with the goal of creating a reusable design.

FEE and Circuit Designs for Streaming Readout at MIT At MIT we will investigate streaming readout starting with small prototype detectors and working with an ASIC chip design and manufacturing company, Alphacore, to develop a range of pre-amplifiers and multi-channel ADC chips that could form the basis for the front-end electronics for a variety of detector types. We would start with simple scintillating fibre or strips with SiPM readout but also develop front-end electronics suitable for GEM detectors and eventually lead tungstate calorimeters working together with our colleagues from INFN Genoa.

The goal of this work would be to establish a catalog of suitable front-end electronics and circuit designs with a known range of parameters that could be used for a variety of detectors. These could be combined with ADC or TDC chips on the same board mounted close to the detector or connected at some distance from the detector with copper or fiber. The interface to signal processing on FPGA boards would also be investigated using simple, low-cost FPGA kits.

In addition to the above, we propose to start the following *new* prototype activities specifically for EIC. **There are funding requests associated with the following.**

Streaming readout for an EIC Calorimeter at INFN Genoa and CUA At INFN-Genova and CUA, we propose to investigate the effect of a streaming readout solution for the electromagnetic calorimeter (EMCAL) of the EIC detector. The EMCAL will be one of the key elements defining the trigger for most, if not all the processes of interest for the realization of the EIC physics program since it will be crucial for tagging the scattered electron. For this reason, we propose the study and development of a streaming data acquisition system optimized for the EMCAL readout. We aim to prototype general solutions that are valid for a fast (O(10-ns)) electromagnetic calorimeter, applicable to any specific technology that will be chosen (crystals, powder/scintillators, Shashlik, ...) and, largely, on the photo-sensors that will be used (SiPMs, APDs, PMTs, ...). The proposed activity includes: Implementation of a 16ch DAQ system based on a commercial FADC board that allows for triggerless readout; Implementation of the full SiPM operation and read-out chain based on discrete components, including: the HV-power supply, the analog front-end to match the requested digitiser input; the signal digitisation and triggerless readout; Development and test of trigger algorithms to be implemented on a on-line farm of CPUs; Performance comparison with results obtained using a traditional FPGA-based trigger DAQ.

The concept will be validated by running preliminary tests on a small array of $PbWO_4$ crystals instrumented with SiPM. The crystals are available and ready to be assembled in a matrix. $PbWO_4$ represents a leading choice for the EIC calorimeters and results of this R&D will be particularly significant to validate the proposed technology in a realistic configuration. Thanks to the high-level programming of the software trigger, it will be easy to adapt the selection criteria to different crystals/configurations for the optimization of the other EM calorimeter options.

Multilayered Architecture for Streaming Readout at Stony Brook We propose to investigate and develop a prototype for a layered architecture for the streaming data transport network. This includes:

1. Study of the most appropriate low-level link protocols (OSI layers 3-5). We want to study the different trade-offs between standard solutions (e.g. IP, TCP, UDP, raw Ethernet frames, continuous serial) with regard to ease of implementation, reliability and overhead.
2. Development of suitable application layer (OSI 7) protocols and solutions for data transport, slow control and configuration management. We want to address architectural questions with regard to the data distribution scheduling, data packaging, etc. This includes FPGA implementations of the core protocols which are then available as public IP for integration into FEE firmwares used by the other sub-projects.
3. Development of a framework prototype for online data processing software, e.g. software-defined event selection, calibration and monitoring.
4. Study of possible architectures for the CPU farm. We want to study in what ways a CPU cluster for online processing could be most effectively retasked to serve for offline analysis or general high performance computing while no data is taken.

This list of projects obviously does not exhaust the parameter space that needs to be investigated for EIC. An aim of the consortium is to provide a framework into which new projects could be plugged in. We expect to attract groups with new interests and projects to join the consortium. We also are aware of the vast knowledge pool associated with many experiments recently converting to or are in the process of converting to streaming-readout type architecture. We also expect to tap this resource, if not though direct collaboration, then through targeted contacts with these groups.

Funding Request covers two areas

- Hardware costs for establishing a few new prototype streaming-readouts.
- The list and definition of relevant parameters will be developed in several intensive targeted workshops by the consortium. Travel funds are requested to support this activity.

BUDGET

The following is the proposed budget for the next year. In case of 20% and 40% reduction, we would reduce the travel funds.

Institute	Equipment	Travel	Sum
Catholic University of America	10k	5k	15k
INFN Genova	10k	5k	15k
Stony Brook	15k	5k	20k
MIT	0	0	0
JLab	0	0	0
Total			50k